

Learning to Match for Multi-criteria Document Relevance

Bilel Moulahi
IRIT, Université Paul Sabatier
118 Route Narbonne,
Toulouse, France
Université de Tunis El Manar
Faculté des Sciences de Tunis
LIPAH, 2092, Tunis, Tunisie
bilel.moulahi@irit.fr

Lynda Tamine
IRIT, Université Paul Sabatier
118 Route Narbonne
Toulouse, France
tamine@irit.fr

Sadok Ben Yahia
Université de Tunis El Manar
Faculté des Sciences de Tunis
LIPAH, 2092, Tunis, Tunisie
sadok.benyahia@fst.rnu.tn

ABSTRACT

In light of the tremendous amount of data produced by social media, a large body of research have revisited the relevance estimation of the users' generated content. Most of the studies have stressed the multidimensional nature of relevance and proved the effectiveness of combining the different criteria that it embodies. Traditional relevance estimates combination methods are often based on linear combination schemes. However, despite being effective, those aggregation mechanisms are not effective in real-life applications since they heavily rely on the non-realistic independence property of the relevance dimensions. In this paper, we propose to tackle this issue through the design of a novel fuzzy-based document ranking model. We also propose an automated methodology to capture the importance of relevance dimensions, as well as information about their interaction. This model, based on the Choquet Integral, allows to optimize the aggregated documents relevance scores using any target information retrieval relevance metric. Experiments within the TREC Microblog task and a social personalized information retrieval task highlighted that our model significantly outperforms a wide range of state-of-the-art aggregation operators, as well as a representative learning to rank methods.

1. INTRODUCTION

A large body of research has focused on the core concept of relevance in information retrieval (IR) [3, 29, 33]. While early work particularly considered the topical relevance, others argued that the user context is highly dependent on many relevance factors that represent the basic clue for relevance assessment [29]. Recently, the increasing availability of user generated content over social media has brought new challenges to multi-relevance estimation because of both the diversity, the task-dependency and the user's context-dependency of the involved relevance features. For instance, consider a user submitting the query "municipal elections" in

the Twitter search system. In such scenario, the user usually looks for very fresh tweets from her/his region rather than other locations, satisfying as much as possible his/her information need (*i.e.*, topical matching between the query and the tweets). While a typical IR system is able to return recent tweets, it often suffers from coverage and ranking problems, as the user may be interested in more complex ranking scenario encompassing other Twitter quality criteria beyond topicality (*e.g.*, authority, credibility and geo-localization dimensions, etc.). Interestingly, studies held through *Searchmetrics*¹ highlighted both the diversity of the features used by modern search engines such as Google+, Twitter and Facebook and the important part of social features particularly.

A wide range of challenging IR applications including mobile IR [19], personalized IR [14, 13, 15, 31] and social IR [24, 16, 21] involve heavily the aggregation of multiple relevance dimensions. In the sake of addressing this challenge, previous solutions are mostly based on classical aggregation functions such as weighted means or linear combination schemes in the form of products and sums. However, these aggregation operators assume that relevance dimensions are *not independent* of each other and generally those including topical relevance are rated as the highest ones in importance [29, 17]. For instance, we show in Figure 1 the distribution of relevance scores of Topic 83: "Stuxnet Worm effects" from the TREC² 2012 Microblog track with respect to recency and topicality criteria. We compute the Spearman Rank correlation coefficient, and we find a notable negative correlation ($\rho = -0.54$), as it may be seen in the lower and higher parts of the graph. Obviously, the performances scores are considerably high either for the topicality criterion or the recency one. Considering the case of such correlation scenarios and bearing in mind that classical aggregation operators are assumed to hold the additive property across individual scores, the latter is unable to account for dependencies between those interacting relevance dimensions. Therefore, using an aggregation mechanism relying on a linear combination scheme can be effective and convenient in some applications, but can also be somewhat inadequate in many IR tasks. Although advanced aggregation operators were recently proposed [13, 18, 17], only a few [17] considered specifically the interactions existing among the relevance dimensions. In practice, this problem is usually avoided by

¹<http://www.searchmetrics.com/en/services/ranking-factors-2013/>

²<http://trec.nist.gov>




Figure 1: Distribution of relevance scores for topic 83, “Stuxnet Worm effects”, of the TREC 2012 Microblog track.

constructing independent criteria [19].

In this paper, we assume a more general scenario where different dependent or independent relevance dimensions are considered within a document retrieval task. We propose a new multidimensional relevance aggregation operator based on the discrete Choquet integral [10]. The latter has been successfully applied on a wide variety of domains during the last decades [20].

More specifically, the contributions of this paper are twofold and mainly include:

- A novel multi-criteria aggregation approach to address the problem of combining relevance estimates and a machine learning driven methodology to train the model;
- A large-scale experimental evaluation within two different IR tasks namely the tweet search task [28] and the personalized social bookmarking task [34].

The remainder of the paper is organized as follows. In Section 2, we briefly survey related work to put our contribution in context. Section 3 describes our multidimensional relevance aggregation operator. In Sections 4 and 5, we describe the experimental setup and then present the experiments and discuss the obtained results. Section 6 concludes the paper and outlines future work.

2. RELATED WORK

In this section, we briefly overview two research directions related to document relevance estimates combination: (i) the relevance concept and the different features it entails; and (ii) relevance criteria combination mechanisms. We present some approaches that are relevant to multidimensional relevance aggregation, and we reveal their significant differences from standard relevance aggregation mechanisms.

2.1 On the Diversity of Relevance Features in IR

The relevance concept in IR has gathered a great attention in the two last decades [3, 29]. Although there is no wide consensus, the general relevance level includes many dimensions such as content, object, validity, situational, affective and belief dimensions [29]. Each dimension refers to a group of criteria considered by the users to make relevance

inferences [29].

The main findings are threefold:

- Relevance dimensions are *not independent* of each others and generally those related to content, which include topical relevance, are rated as the highest ones in importance, but *interact* with other dimensions [29, 17];
- A few and finite relevance dimensions are considered jointly by the users to assess relevance;
- The importance of the dimensions depends on tasks in progress and class of users.

Considering these findings and tasks specificities, works in many recent IR applications such as mobile IR [11, 19], social IR [24, 16, 27] and personalized IR [31], attempt to go beyond the classical content-matching dimension.

For instance, in mobile IR, relevance is based on the user’s situation (location) and social (surrounding persons) dimensions [19, 11]. In social networking services like Twitter, a variety of micro-blogging specific criteria such as topic matching and authority belonging respectively to content and validity dimensions, are investigated to compute relevance [27]. In the same line of research, authors in [16] propose to rank tweets according to the number of followers, the number of mentions as well as the authority of the microblogger, computed by applying PageRank algorithm on retweet social network. The content relevance feature is always the most important feature for all search settings. The authors in [21] proposed new social features for social search in micro-blogging networks such as Twitter. They proposed new relevance factors such as the social importance of microbloggers and the temporal magnitude of tweets. The social importance is considered as an indicator of tweets credibility, and refers to the influence of the microblogger on the social network. The temporal magnitude of microblogs is estimated based on temporal neighbors that present similar query terms.

In [9], Chelaru et al. investigated the impact of social features on the effectiveness of video retrieval in Youtube. Social features including likes, dislikes, comments, favorites, etc., are those generated from implicit or explicit interaction with the system. The authors showed that social features are valuable and have a great potential to improve the video retrieval performance when combined to a basic feature like topicality.

In social tagging systems, there are also various indicators of relevance representing the users and tags that could be combined in different models. In a tag recommendation oriented problem, Belém et al. [1] jointly exploited three dimensions of relevance: (i) co-occurrence of terms with tags previously assigned to the object; (ii) terms extracted from multiple textual features, such as title and description; and (iii) relevance metrics such as “Term Frequency”. In a more recent work [2], the authors considered both aspects “Novelty” and “Diversity” to address the problem of tag recommendation. Novel tags are those that are observed very often in the application, and the diversity of a list of recommended tags is estimated by the average semantic distance between each pair of tags in the list. The authors claimed that both aspects are distinct but related concepts. The proposed ranking model combining those metrics has been tested on 3

popular social media websites: LastFM, Youtube and YahooVideo.

2.2 On Relevance Aggregation in IR

Regarding the second line of research that addresses combination of multiple relevance estimates, we identify in the sequel three main strategies.

Linear combination approaches. Researches involving evidence combination from multiple relevance features are often based on linear combination functions [31, 27] due to their simplicity and their relative effectiveness. For instance, in the TREC Microblog tracks [32], combination of the different relevance indicators is mostly based on linear functions in the form of products or sums [24]. Gerani *et al.* [18] have proposed a multi-criteria aggregation model allowing to generate a global score that does not necessarily require the comparability of the combinable individual scores. The authors rely on the Alternating Conditional Expectation Algorithm and the BoxCox model to analyze the incomparability problem and perform a score transformation whenever necessary. More recently, Eickhoff *et al.* [17] introduced a more advanced aggregation mechanism which consists in a statistical framework based on *copulas*. This model addresses the multidimensional relevance assessment, where documents are described by several correlated relevance criteria. The authors compare the copula model to three combination schemes: Sum, Product and weighted linear combination of relevance scores. They argue that these baselines assume independence across all criteria and can be expected to be too naive in some settings where dependence is given [17]. The presented results show that copulas models outperform linear combination mechanisms in two IR tasks namely: opinionated blogs and personalized bookmarks; however, results in another IR application, *i.e.*, retrieval of child-friendly websites, show that linear combination scheme outperforms the copulas models.

Another family extensively studied in the literature, is that of ordered weighted averaging functions (Owa) initiated by Yager [35]. In [4], the authors proposed an aggregation scheme that adopt the same idea of the Owa [35] operator and uses a weighting method that gives more importance to the terms with high relevance degree to minimize the impact of terms having low scores on the global final evaluation.

Prioritized aggregation operators. These aggregation operators model a priority relationship over the set of criteria. Yager [36] introduced a prioritized scoring operator and a closely related prioritized averaging operator that makes the weights associated to each criterion dependent upon the satisfaction of the higher preferred criterion. In the IR field, Celia *et al.* [13] proposed a multidimensional representation of relevance through a general prioritized aggregation scheme involving two operators namely, “And” and “Scoring” [12]. While the Scoring operator is inspired by the prioritized operator proposed in [36], the “And” operator is based on a refinement of the $\$minT$ operator. The peculiarity of these operators consists in the fact that if a criterion is not important for the user, its value does not affect the overall performance score, *i.e.*, the weight of a less important criterion should be proportional to the satisfaction degree of more important criteria. The aggregation operators are evaluated in a personalized IR setting. The latter show notable performance improvements when compared to the average operator.

Learning to Rank Approaches. Instead of proposing carefully designed ranking models based on heuristics principles, a recent emerging thread of research is to apply machine learning algorithms in order to combine multiple relevance features [6, 7, 23]. Given a training set of queries and the associated ground truth containing document labels (relevant, irrelevant), the objective is to optimize the relevance retrieval metrics with the goal of improving the overall search result quality. Optimization is formulated as learning a ranking function in order to minimize a loss function in the training data (eg., number of miss-ordered document pairs). Each query-document pair is represented by a feature vector consisting of different variables which are functions of the content of the document. In the tweet search task, different studies [16, 24, 8] propose learning to rank approaches in order to combine several types of twitter features. These work show the benefits of considering machine learning methods in aggregating relevance criteria. Algorithms like RankSVM and decision tree based learning to rank approaches often show improvements when used in IR tasks [23].

However, despite being effective in many IR applications [23], these methods tend to offer only limited insight on how to consider importance and interaction between the groups of features that are mapped to different relevance dimensions [17].

Unlikely, we propose to investigate the combination of general level relevance dimensions using a fuzzy-based aggregation operator. More oriented to the specific problem of relevance aggregation, our method is able to address the property of interaction between dimensions through an integral aggregation operator, namely the Choquet integral, *w.r.t* a fuzzy measure expressing both their individual and joint importance. These properties appear to be appealing from an IR perspective. In our work [26], we have successfully proposed a general personalized approach based on Choquet in a contextual suggestion IR task.

3. COMBINING RELEVANCE ESTIMATES WITH THE CHOQUET INTEGRAL

3.1 Choquet-based Relevance Aggregation: Presentation

An aggregation operator is a function that maps several inputs from a given interval (*e.g.*, $[0 \dots 1]$) to a single output in the same interval. In a typical IR setting, inputs consists in performance scores (RSV³) obtained *w.r.t* each relevance dimension. In fact, the difficulty in the multidimensional relevance aggregation problem is twofold:

- *Estimation of the relevance criteria importance:* identifying which individual criterion and/or subset of criteria need to be enhanced *vs.* weakened regarding the IR task at hand;
- *Aggregation and document ranking:* accurately combining the relevance criteria by taking into account their dependency. Based on the documents global scores obtained by combining the partial performance scores *w.r.t* each criterion, it is decided whether a document should be ranked better than another in a ranking.

³Retrieval Status Value.

Consider an IR scenario where $\mathcal{D} = \{d_1, d_2, \dots, d_M\}$ is the set of documents, $\mathcal{C} = \{c_1, c_2, \dots, c_N\}$ is the set of relevance criteria and q a given query. Let $RSV_{c_i}(q, d_j)$ be the performance scores of document $d_j \in \mathcal{D}$, obtained w.r.t the relevance criterion $c_i \in \mathcal{C}$. The task of combining the performance scores of d_j w.r.t all $c_i \in \mathcal{C}$ is called *aggregation*. Formally, the general aggregation function, denoted by $\mathcal{F} : \mathbb{R}^N \rightarrow \mathbb{R}$, that computes the global score of document d_j in response to query q , is defined as follows.

$$\begin{aligned} & RSV_{c_1}(q, d_j) \times \dots \times RSV_{c_N}(q, d_j) \\ & \rightarrow \mathcal{F}(RSV_{c_1}(q, d_j) \\ & \quad \dots, RSV_{c_N}(q, d_j)) \end{aligned} \quad (1)$$

Where $RSV_{c_i}(q, d_j)$ is the performance score of d_j w.r.t criterion c_i . As we rely on the Choquet operator, to each subset of criteria is associated a fuzzy measure, also called *capacity*, that reflects its importance.

DEFINITION 1. Let I_C be the set of all possible subsets of criteria from \mathcal{C} . A fuzzy measure is a normalized and a monotone function μ from I_C to $[0 \dots 1]$ such that:

$\forall I_{C_1}, I_{C_2} \in I_C$, if $(I_{C_1} \subseteq I_{C_2})$ then $\mu(I_{C_1}) \leq \mu(I_{C_2})$, with $\mu(I_\emptyset) = 0$ and $\mu(I_C) = 1$.

For the sake of notational simplicity, $\mu(I_{C_i})$ will be denoted by μ_{C_i} . The value of μ_{C_1} can be interpreted as the importance degree of the interaction between the criteria involved in the subset C_1 .

The Choquet integral based-relevance aggregation function built on such a fuzzy measure is defined as follows.

DEFINITION 2.

$$\begin{aligned} & Ch_\mu(RSV_{c_1}(q_r, d_j), \dots, RSV_{c_N}(q_r, d_j)) \\ & = \sum_{i=1}^N \mu_{c_i, \dots, c_N} \cdot (rsv_{(i)j} - rsv_{(i-1)j}) \end{aligned} \quad (2)$$

Where Ch_μ is the Choquet aggregation function, $rsv_{(i)j}$ is the i -th element of the permutation of $RSV(q_r, d_j)$ on criterion c_i , such that $(0 \leq rsv_{(1)j} \leq \dots \leq rsv_{(N)j})$.

The value of $\mu_{I_{C_i}}$ can be interpreted as the importance degree of the interaction between the criteria involved in subset C_i . Note that if μ is an additive measure, the Choquet integral corresponds to the weighted mean. Otherwise, it requires fewer than $(2^N - 1)$ capacity measures in the case where the fuzzy measure is k -order additive, i.e., $\mu_A = 0$ for all criteria subsets $A \subseteq \mathcal{C}$ with $|A| > k$.

As previously stated, the Choquet operator exhibits a number of properties that appear to be appealing from an IR point of view. From a theoretical perspective, since it is built on the concept of fuzzy measures, it allows modeling flexible interactions and considering complex dependencies among criteria [20]. To facilitate the task of interpreting the Choquet integral behavior, we exploit two parameters namely, the ‘‘importance indice’’ and the ‘‘interaction indice’’ [20] that offer readable interpretations and qualitative understanding of the resulting aggregation model. For more details on the interaction and importance indices, see our extended work in [25].

DEFINITION 3. Importance index: Let μ_{c_i} be the weight of relevance criterion c_i and $\mu_{C_r \cup c_i}$ its marginal contribution

to each subset $Cr \in \mathcal{C}$ of other criteria. The importance index [30] of c_i w.r.t a fuzzy measure μ is then defined as the mean of all these contributions:

$$\phi_\mu(c_i) = \sum_{Cr \subseteq \mathcal{C} \setminus \{c_i\}} \frac{(N - |Cr| - 1)! \cdot |Cr|!}{N!} [\mu_{Cr \cup c_i} - \mu_{Cr}] \quad (3)$$

$\phi_\mu(c_i)$ measures the average contribution that criterion (c_i) brings to all the possible combinations of criteria.

By introducing this indice, the overall importance of criterion c_i is no longer solely determined by its weight μ_{c_i} but also by its contribution to each subset of other criteria.

DEFINITION 4. Interaction index: Let $(\Delta_{c_i c_j} \mu_{Cr})$, with $Cr = \mathcal{C} \setminus \{c_i, c_j\}$, be the difference between the marginal contribution of criterion c_j to every combination of criteria that contains criterion c_i , and a combination from which criterion c_i is excluded:

$$(\Delta_{c_i c_j} \mu_{Cr}) = [\mu_{(\{c_i c_j\} \cup Cr)} - \mu_{(c_i \cup Cr)}] - [\mu_{(c_j \cup Cr)} - \mu_{Cr}] \quad (4)$$

This expression is defined to appraise the strength among two criteria c_i and c_j . When this latter expression is positive (resp. negative) for any $Cr \in \mathcal{C} \setminus \{c_i, c_j\}$, we say that both criteria c_i and c_j positively (resp. negatively) interact (i.e., the contribution of criterion c_j is higher with the presence of criterion c_i).

The interaction index among two measures is thus defined as follows:

$$I_\mu(c_i, c_j) = \sum_{Cr \subseteq \mathcal{C} \setminus \{c_i, c_j\}} \frac{(N - |Cr| - 2)! \cdot |Cr|!}{(N - 1)!} (\Delta_{c_i c_j} \mu_{Cr}) \quad (5)$$

The interaction value, which falls into the interval $[-1..1]$, is zero when both criteria are independent and it is positive (resp. negative) whenever the interaction between them is positive (resp. negative).

3.2 Training the fuzzy measures within an IR task

The objective of the training step here is to optimize the fuzzy measures w.r.t a target IR measure (e.g. $P@X$) by identifying the values of the Choquet capacities to be used in the aggregation process. Figure 2 illustrates the adopted methodology. As in machine learning algorithms, the typical training data required for learning the Choquet fuzzy measures includes a set of training queries q_k and for each query, a list of ranked documents d_{kj} represented by pre-computed vectors containing performance scores $RSV_{c_i}(q_k, d_{kj})$. $RSV_{c_i}(q_k, d_{kj})$ denotes the score of d_{kj} w.r.t criterion c_i in response to query q_k . Each document is annotated with a rank label l_{kj} (e.g., relevant or irrelevant). As presented in Figure 2, the training involves two main steps, described in what follows.

Tuning: the objective of this step is to generate and tune the capacity values in order to identify the best initial capacity values. In a first stage, we start by generating and initializing a set of capacity combinations, where each capacity combination $\mu^{(\cdot)}$ represents the set of capacity values

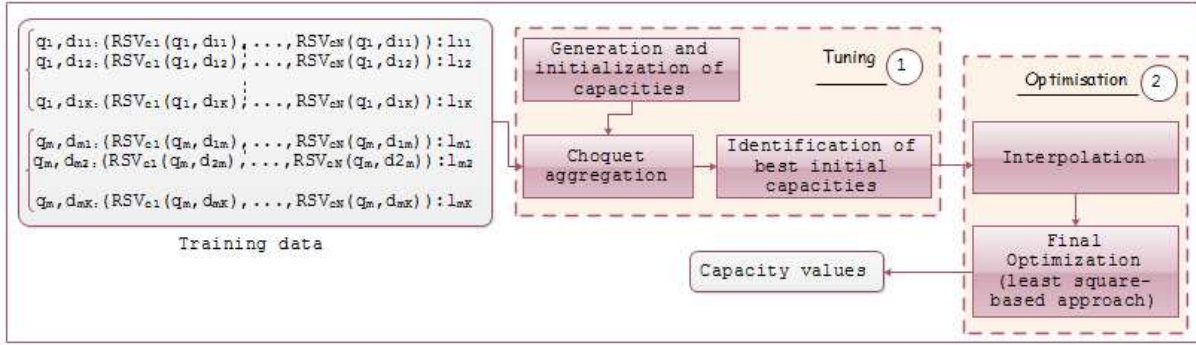


Figure 2: General Paradigm of Training the Fuzzy Measures (Capacities) for all Criteria and Subsets of Criteria

assigned to each criterion and subset of criteria. In the case of N relevance criteria, each capacity combination includes $(2^N - 1)$ capacity values. For instance, if we consider three relevance criteria, a capacity combination involves the values $(\{\mu_{c_1}; \mu_{c_2}; \mu_{c_3}; \mu_{c_1, c_2}; \mu_{c_1, c_3}; \mu_{c_2, c_3}\})$. Those values fall into $[0 \dots 1]$ and are tuned for all criteria, such that their sum is equal to 1. The value are computed with a step equal to 0.1. In the case of three criteria, we fix the capacity value of a criterion (e.g., 0.1) and we tune the remaining values (e.g., from 0.8 to 0.1 for the first; and from 0.1 to 0.8 for the second). the process is repeated for each criterion. The tuning is conceivable since there is generally a few relevance dimensions [29]. However, when the number of criteria is strictly higher than 3, we can avoid the tuning complexity by relying on sub-families of capacities namely 2-additive measures [20], requiring less coefficients to be defined and assuming that there is no initial interaction among subsets of more than 2 criteria. In a second stage, we used the Choquet integral operator to aggregate the different relevance criteria estimates of the performance scores in the training data *w.r.t* to each generated capacity combination $\mu^{(\cdot)}$. Then, we select the best capacity combination $\mu^{(*)}$ giving the best results *w.r.t* a target IR measure (e.g., $P@X$).

Optimisation: the objective of this step is to optimize the initial capacity combination $\mu^{(*)}$ obtained in the previous step. We start by $\mu^{(*)}$, we pull the top K documents returned by each training query q_k and we interpolate the documents scores to boil down the non relevant ones. Interpolation is based on performance scores of the training documents based on $\mu^{(*)}$. The underlying objective is to affect relevant documents higher scores than non relevant ones. Finally, we proceed to the application of the Least-squares based optimization. Obtained capacity values are used within each relevance criterion and subset of relevance criteria.

4. EXPERIMENTAL DESIGN

Experimental evaluation is based on two IR evaluation frameworks, namely the tweet search task [28] and the personalized social bookmarking IR task [34]. This section describes these tasks, the used data, baselines and metrics.

4.1 Tasks

4.1.1 Tweet Search Task.

We exploit the dataset and topics of the TREC 2011 and 2012 Microblog tracks. The datasets include more than 16

million tweets and more than 5 million users [28]. TREC Microblog 2011 track includes (49) topics, used for training the capacity values and TREC Microblog 2012 includes (60) topics used for testing. We make use of three important and widely used relevance criteria [27]: *topicality*, *recency* and *authority*. To deal with the topicality relevance criterion, we propose to use the Okapi BM25 ranking, then, we introduce the authority and the recency relevance dimensions as query independent measures, as done in [27]. The analysis of the criteria importance using the importance indice (Cf., Section 3.1) [20] reveals a high importance of topicality with a value of 0.631. The recency relevance criterion is also given a quite high importance (of about 0.25) compared to the authority relevance dimension (0.12). This is not surprising as far as a user usually seeks for topically relevant documents rather than those which are authoritative or even more recent.






Figure 3: Criteria Interaction Indice.

Figure 3 shows the values of the interaction indice between topicality, recency and authority, denoted respectively by *To*, *Re* and *Au* within both TREC Microblog 2011 and 2012 track topics (referred respectively as “TMTOPICS 2011” and “TMTOPICS 2012”). From this Figure, we can see that the authority criterion does not bring any contribution when it is combined with topical relevance criteria. Moreover, we notice a positive interaction between the *topicality* and *recency* relevance criteria. This explains the higher contribution of these two criteria on the overall global scoring when they are present together, and allows us to evaluate our approach under the presence of interdependent relevance criteria.

4.1.2 Personalized social bookmarking task.



(a) Capacities values effectiveness within the tweet search task. The x -axis represents some of the 21 trained capacities combinations.



(b) Capacities values effectiveness within the social personalized task.

Figure 4: Capacities tuning within both datasets.

With the advent of the Web 2.0, social tagging systems, such as *e.g.*, *Del.icio.us*⁴ and *Flickr*⁵, have exponentially grown both in terms of users and contents. We exploit a compiled collection of about 33k delicious bookmarks and 12k tags.⁶ The dataset contains evaluation information from 35 users over 177 search topics [34] according to two relevance dimensions namely, the topical relevance (To) of bookmarks given a topic and its personal or user relevance (Us) given a user. In our evaluation protocol, we use 75% of the topics for learning the capacity values and we exploit the remaining topics for the testing phase.

We compute the importance indices and we found that both of them have quite similar importance with indices of about 0.48 for the topical relevance and 0.51 for the user one. The interaction value between the considered relevance criteria ($To - Us$) is about 0.028, which is quite low to argue that they are really independent.

4.2 Baseline runs and metrics

Firstly, we compare our approach to classical aggregation operators such as the linear combination scheme (LCS) and the OWA [35], as well as the two competing prioritized operators namely, AND and SCORING [13]. Then, we compare

⁴<http://www.delicious.com>

⁵<http://www.flickr.com>

⁶<http://ir.ii.uam.es/~david/webdivers/>

it to state-of-the-art learning to rank algorithms namely RANKSVM [22], RANDOM FOREST [5] and λ -MART. With respect to each task specificities and guidelines, we used: 1) for the tweet search task, the $P@5$, $P@10$, $P@20$, $P@30$, $P@100$ and MAP measures. We emphasize that the official measure of the track is $P@30$; 2) for the personalized bookmarking task, given that only the top 5 relevant results are provided with each topic, we make use of the $P@5$ measure as recommended in [34]. In these experiments, significance testing is based on the t-student statistic for both datasets.

5. RESULTS AND DISCUSSION

5.1 Training the fuzzy measures

To identify the Choquet capacity values, we adopt the methodology illustrated in Figure 2 (Cf., Section 3.2). The value of the top retrieved documents K is set to 100. For the tweet search task, we instantiate Q_{learn} and $qrels$ by the TREC Microblog 2011 topics and the corresponding relevance assessments respectively and use the official measure $P@30$ for the tuning. For the personalized social bookmarking task, we used random 75% of the topics for the training phase and use the $P@5$ measure for the tuning. The trained combination capacities in this learning phase, will be denoted by $\mu^{(i)}$. The initial capacity value of each criterion is obtained with a step of 0.1 such that the sum of the three capacities is 1, leading to 21 and 9 capacity values respectively

Operator	Precision					MAP	% change	
	P@5	P@10	P@20	P@30	P@100			
LCS	0.1965	0.1860	0.1833	0.1854	0.1309	0.0928	+20,73%	§
OWA	0.1828	0.1879	0.1767	0.1764	0.1248	0.0882	+23,73%	★
AND	0.1828	0.1793	0.1767	0.1764	0.1233	0.0882	+23,73%	★
SCORING	0.2000	0.2018	0.1982	0.1977	0.1475	0.1091	+14,52%	★
RANKSVM-L	0.2207	0.2276	0.2207	0.2213	0.1586	0.1213	+4,32%	★
RANKSVM-RBF	0.0966	0.0828	0.0733	0.0856	0.0993	0.0665	+62,99%	★
RF	0.1000	0.0810	0.0681	0.0687	0.0628	0.0464	+70,68%	★
λ-MART	0.2931	0.2276	0.2092	0.2043	0.1856	0.1321	+11,67%	★
CiFA	0.2379	0.2362	0.2422	0.2313	0.1614	0.1295	—	
	+7.22%	+3.64%	+8.87%	+4.32%	+1.73%	+6.33%		★

Table 1: Comparative evaluation of retrieval effectiveness. % change indicates the CiFA improvements in terms of $P@30$. The symbols § and ★ denote the student test significance: "§": $0.05 < t \leq 0.1$; "★": $t \leq 0.01$. The last row shows the CiFA improvement in terms of $P@X$ and MAP with the best baseline (*i.e.*, RANKSVM).

	LCS	OWA	AND	SCORING	RankSVM	CiFA
P@5	0.6310	0.6310	0.6286	0.6310	0.6286	0.6310
% change	0%	0%	+0,003%	0%	+0,003%	—
	***	***	***	***	***	

Table 2: Comparative evaluation of retrieval effectiveness within the personalized social bookmarking task. The symbol ★ denotes the student test significance: "★★★": $t \leq 0.01$.

for the tweet search task and the personalized bookmarking task. Then, we select the combination ($\mu^{(*)}$) achieving the highest average value of the precision measure. Figure (3a) and (3b) give an overview of the training results for the tweet search dataset and the social bookmarking dataset respectively; the optimal combination $\mu^{(*)}$ is highlighted in black.

From Figure (3a) we can see that, as expected from the importance and interaction analysis presented in Sections 4.1.1 and 4.1.2, the tweet search task is more sensitive to the topical criterion To than the authority Au and recency criteria Re , including values of about 0.8 for μ_{To} and 0.1 for μ_{Au} and μ_{Re} , respectively. After the application of the Least squares based optimization method, we obtained $\mu^{(**)}$ which is composed by: ($\mu_{To} = 0.705$, $\mu_{Re} = 0.215$, $\mu_{Au} = 0.025$, $\mu_{\{To, Re\}} = 0.973$, $\mu_{\{To, Au\}} = -0.14$, $\mu_{\{Re, Au\}} = -0.25$). We notice that the capacity values obtained for $\mu^{(**)}$ on the subsets $\{To, Au\}$ and $\{Re, Au\}$ are negative ones. This suggests that the authority criterion does not appear to be a good factor when combined with topicality or recency, and this explains the negative capacities assigned to $\mu_{\{To, Au\}}$ and $\mu_{\{Re, Au\}}$, which corroborates previous results [27].

From Figure (3b) we observe that the $P@5$ values *w.r.t* to the different capacity combination values are quite close. It comes somewhat unexpected, that the $P@5$ is either high when the user relevance (Us) is quite more important ($\mu^{(7)}$) or even when the topical relevance (To) is highly valued ($\mu^{(2)}, \mu^{(3)}$). As the $P@5$ is reached for ($\mu^{(7)}$) and the latter yields more importance to the user relevance, we select it as an initial capacity for identifying the fuzzy measures with the Least squares based-optimization method. The capacity

values returned are quite similar to those obtained by the importance indice in Section 4.1.2, giving a capacity (μ_{Us}) of about 0.5 to the user relevance, 0.47 for the topical one (μ_{To}) and 0.028 for the subset of both criteria ($\mu_{\{Us, To\}}$).

5.2 Measuring the Retrieval Effectiveness

This subsection presents the results of the testing phase. For this aim, we use the 60 topics from the TREC Microblog 2012 track, and the 25% remaining topics of the personalized bookmarking task. The baseline parameters have been tuned using the same learning datasets used for the choquet based operator. For the learning to rank method, we used the open source code of RANKSVM from [22].

5.2.1 Relevance Estimation Within the Tweet Search Task

Table 1 reports the retrieval performance obtained by our choquet based aggregation operator, denoted CiFA in the remainder, in comparison with the baselines. Note that we ran a series of experiments to optimize the learning to rank methods parameters. Those experiments are run with five cross validation within the same learning set used to find the Choquet capacities values. For RANKSVM we tested the linear (RANKSVM-L model and the Radial Basis Function (RANKSVM-RBF) kernel. Then, we have tried different values of the parameter C ($10^{-5}..10^2$) for the former, and various pairs of (C , σ) for the latter. Through cross validation, we found that $C = 0.00005$ gives the best performance for the linear RANKSVM-L and found that the pair (8.88,7.77) is the best setting for the RBF kernel. For the tree-based learning to rank methods, we found that they are

quite insensitive to parametrization. Using cross validation, we found that a number of trees of about 1000, a learning rate equal to 0.1 and a number of leaves equal to 2 gives the best performance for λ -MART.

From Table 1, we can see that CiFA overpasses all the baselines. We notice that the performance improvements are more important for the classical aggregation operators; they reach 23,73% compared to the OWA operator and 20,73% compared to the LCS operator. For the SCORING operator, the significant improvement is less important. As we considered the prioritization scenario $Sc_1: \{topicality\} \succ \{recency\} \succ \{authority\}$, giving the best $P@30$ average over the other possible prioritization scenarios, we can conclude that the obtained difference of performance, in favor of CiFA, is explained by the consideration of the interactions existing among the set of criteria, that we involved by means of the fuzzy measures. Thus, the global scores can no longer be biased by dependent relevance criteria or overestimated by those highly scored than the other ones. For the AND operator, the improvement difference is sharply better. The obtained results are likely due to the fact that it is mainly based on the MIN operator, which could penalize tweets highly satisfied by the least important criteria. Roughly speaking, if there are many tweets highly scored *w.r.t* the authority criterion (which is likely the case), its overall satisfaction degrees would be biased by this relevance criterion. Last but not least, we can see that CiFA significantly outperforms the linear RANKSVM which represents here the best baseline. The improvement varies between 4.32% to 8.87%, and it is quite high for the $P@5$, $P@10$ and $P@20$, enhancing thus the quality of the first retrieved tweets. Improvements for the other learning to rank algorithms is more important. They are higher for RANKSVM-RBF and RF but less important for λ -MART, with an improvement of about 11,67%.

5.2.2 Relevance estimation within the social personalized IR task.

As we can see from Table 2, the performance of CiFA as well as the baselines are quite close. Apart from the evaluation measure used, this is likely due to the number of criteria involved as well as the slight independence existing between them, as shown through the interaction analysis (Cf. Subsection 4.1.2). Therefore, the behavior of the Choquet integral is quite similar to that of a linear combination scheme which explains the results reported in Table 2. Interestingly, these results also highlight the fact that the Choquet operator performance is stable despite criteria independence and leads to comparable results than the baseline operator results at least when the number of criteria is quite small.

6. CONCLUSION AND FUTURE WORK

We presented a novel general multi-criteria framework for multidimensional relevance aggregation. Our approach relies on a fuzzy integral method based on the well studied and theoretically justified Choquet mathematical operator. The proposed operator supports the observation that relevance dimensions, measurable through criteria, may interact and have different weights (importance) according to the task at hand. Besides, the approach allows analyzing the resulting model criteria behavior with a readable interpretation through the Shapley and interaction indices. A set of experiments were conducted on TREC Microblog datasets as

well as a personalized social bookmarking corpus and showed that our operator significantly outperforms both supervised and unsupervised aggregation methods. The main limitation of this work concerns the computational complexity of the learning algorithm in the case where the number of relevance dimensions is high. This specific issue can be addressed by considering other properties of the Choquet integral to be drawn within the document rankings.

In future, we plan to investigate how to integrate the user or class of users (eg. children, students) as a dependent variable within the framework in order to design personalized aggregated rankings. It is possible to achieve this by learning a user-based fuzzy measure and assume independence between user classes.

7. REFERENCES

- [1] F. Belém, E. Martins, T. Pontes, J. Almeida, and M. Gonçalves. Associative tag recommendation exploiting multiple textual features. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '11, pages 1033–1042, New York, NY, USA, 2011. ACM.
- [2] F. M. Belém, E. F. Martins, J. M. Almeida, and M. A. Gonçalves. Exploiting relevance, novelty and diversity in tag recommendation. In *Proceedings of the 18th Brazilian Symposium on Multimedia and the Web*, WebMedia '12, pages 297–300, New York, NY, USA, 2012. ACM.
- [3] P. Borlund. The concept of relevance in IR. *Journal of the American Society for Information Science and Technology*, 54(10):913–925, 2003.
- [4] M. Boughanem, Y. Loiseau, and H. Prade. Rank-ordering documents according to their relevance in information retrieval using refinements of ordered-weighted aggregations. In *Proceedings of the Third international conference on Adaptive Multimedia Retrieval: user, context, and feedback*, AMR'05, pages 44–54, Berlin, Heidelberg, 2006. Springer-Verlag.
- [5] L. Breiman. Random forests. *Mach. Learn.*, 45(1):5–32, 2001.
- [6] C. Burges, T. Shaked, E. Renshaw, A. Lazier, M. Deeds, N. Hamilton, and G. Hullender. Learning to rank using gradient descent. In *Proceedings of the 22nd international conference on Machine learning*, ICML '05, pages 89–96, New York, NY, USA, 2005. ACM.
- [7] Z. Cao, T. Qin, T.-Y. Liu, M.-F. Tsai, and H. Li. Learning to rank: from pairwise approach to listwise approach. In *Proceedings of the 24th international conference on Machine learning*, ICML '07, pages 129–136, New York, NY, USA, 2007. ACM.
- [8] Y. Chang, A. Dong, P. Kolari, R. Zhang, Y. Inagaki, F. Diaz, H. Zha, and Y. Liu. Improving recency ranking using twitter data. *ACM Trans. Intell. Syst. Technol.*, pages 4:1–4:24, 2013.
- [9] S. Chelaru, C. Orellana-Rodriguez, and I. S. Altingövd. Can social features help learning to rank youtube videos? In *Proceedings of the 13th International Conference on Web Information Systems Engineering*, volume 4 of *WISE'12*, pages 552–566, 2012.
- [10] G. Choquet. Theory of capacities. *Annales de l'Institut Fourier*, 5:131–295, 1953.

- [11] G. Cong, C. S. Jensen, and D. Wu. Efficient retrieval of the top-k most relevant spatial web objects. *Proc. VLDB Endow.*, 2:337–348, 2009.
- [12] C. da Costa Pereira, M. Dragoni, and G. Pasi. Multidimensional relevance: A new aggregation criterion. In *Proceedings of the 30th European conference on Advances in Information Retrieval*, ECIR’ 09, pages 264–275, 2009.
- [13] C. da Costa Pereira, M. Dragoni, and G. Pasi. Multidimensional relevance: Prioritized aggregation in a personalized information retrieval setting. *Inf. Process. Manage.*, 48(2):340–357, 2012.
- [14] M. Daoud, L. Tamine, and M. Boughanem. A personalized graph-based document ranking model using a semantic user profile. In *UMAP*, pages 171–182, 2010.
- [15] M. Daoud, L. Tamine, and M. Boughanem. A personalized search using a semantic distance measure in a graph-based ranking model. *J. Information Science*, 37(6):614–636, 2011.
- [16] Y. Duan, L. Jiang, T. Qin, M. Zhou, and H.-Y. Shum. An empirical study on learning to rank of tweets. In *Proceedings of the 23rd International Conference on Computational Linguistics*, COLING ’10, pages 295–303, Stroudsburg, PA, USA, 2010. Association for Computational Linguistics.
- [17] C. Eickhoff, A. P. de Vries, and K. Collins-Thompson. Copulas for information retrieval. In *Proceedings of the 36th annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Dublin, Ireland, 2013. ACM.
- [18] S. Gerani, C. Zhai, and F. Crestani. Score transformation in linear combination for multi-criteria relevance ranking. In *Proceedings of the 34th European conference on Advances in Information Retrieval*, ECIR’12, pages 256–267, Berlin, Heidelberg, 2012. Springer-Verlag.
- [19] A. Göker and H. Myrhaug. Evaluation of a mobile information system in context. *Inf. Process. Manage.*, 44(1):39–65, 2008.
- [20] M. Grabisch, T. Murofushi, M. Sugeno, and J. Kacprzyk. *Fuzzy Measures and Integrals. Theory and Applications*. Physica Verlag, Berlin, 2000.
- [21] L. B. Jabeur, L. Tamine, and M. Boughanem. Uprising microblogs: A bayesian network retrieval model for tweet search. In *Proceedings of the 27th Annual ACM Symposium on Applied Computing*, SAC ’12, pages 943–948, New York, NY, USA, 2012. ACM.
- [22] T. Joachims. Training linear svms in linear time. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, KDD ’06, pages 217–226, New York, NY, USA, 2006. ACM.
- [23] H. Li. *Learning to Rank for Information Retrieval and Natural Language Processing*. Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers, 2011.
- [24] D. Metzler and C. Cai. Usc/isi at trec 2011: Microblog track. In *Proceedings of The Twentieth Text REtrieval Conference*, 2011.
- [25] B. Moulahi, L. Tamine, and S. B. Yahia. iAggregator: Multidimensional relevance aggregation based on a fuzzy operator. *Journal of the Association for Information Science and Technology*, 65(10):2062–2083, 2014.
- [26] B. Moulahi, L. Tamine, and S. B. Yahia. Toward a personalized approach for combining document relevance estimates. In *User Modeling, Adaptation, and Personalization - 22nd International Conference UMAP*, volume 8538. Springer, 2014.
- [27] R. Nagmoti, A. Teredesai, and M. De Cock. Ranking approaches for microblog search. In *Proceedings of the 2010 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology - Volume 01*, WI-IAT ’10, pages 153–157, Washington, DC, USA, 2010. IEEE Computer Society.
- [28] I. Ounis, C. Macdonald, J. Lin, and I. Soboroff. Overview of the trec-2011 microblog track. In *Proceedings of the 20th Text REtrieval Conference*, 2011.
- [29] T. Saracevic. Relevance: A review of the literature and a framework for thinking on the notion in information science. part iii: Behavior and effects of relevance. *Journal of the American Society for Information Science*, 58(13):2126–2144, 2007.
- [30] L. S. Shapley. A value for n-person games. In H. W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games, Volume II*, volume 28 of *Annals of Mathematics Studies*, pages 307–317, Princeton, 1953. Princeton University Press.
- [31] A. Sieg, B. Mobasher, and R. Burke. Web search personalization with ontological user profiles. In *Proceedings of the sixteenth ACM Conference on Information and Knowledge Management*, CIKM ’07, pages 525–534, New York, NY, USA, 2007. ACM.
- [32] I. Soboroff, I. Ounis, C. Macdonald, and J. Lin. Overview of the trec-2012 microblog track. In *Proceedings of the 21th Text REtrieval Conference*, 2012.
- [33] A. R. Taylor. User relevance criteria choices and the information search process. *Inf. Process. Manage.*, 48(1):136–153, 2012.
- [34] D. Vallet and P. Castells. Personalized diversification of search results. In *Proceedings of the 35th annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 841–850. ACM, 2012.
- [35] R. R. Yager. On ordered weighted averaging aggregation operators in multicriteria decision making. *IEEE Transactions On Systems Man And Cybernetics*, 18(1):183–190, 1988.
- [36] R. R. Yager. Prioritized aggregation operators. *International Journal of Approximate Reasoning*, 48(1):263–274, 2008.